

CME 323: Distributed Algorithms and Optimization

Instructor: Reza Zadeh (rezab@stanford.edu)

TA: Anuj Nagpal (anujnag@stanford.edu)

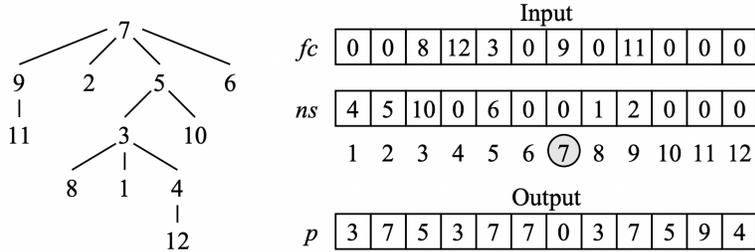
HW#2 – Due Thursday April 28th (on Gradescope)

Total Score – 80 points

1. **(10 points) List Prefix Sums:** List Prefix Sums is the task of determining the sum of all the elements before each element in a list. Let us consider the following simple variation.
 - Select each element from the list randomly and independently with probability $1/\log n$ and add it to a set S . Add the head of the list to this set, and mark all these elements in the list.
 - Start from each element $s \in S$, and in parallel traverse the lists until you find the next element in S (by detecting the mark) or the end of the list. For $s \in S$, call this element found in this way $\text{next}(s)$. While traversing, calculate the sum from s to $\text{next}(s)$ (inclusive of s but exclusive of $\text{next}(s)$), and call this $\text{sum}(s)$.
 - Create a list by linking each $s \in S$ to $\text{next}(s)$ and with each node having weight $\text{sum}(s)$.
 - Compute the List Prefix Sums on this list using pointer jumping. Call the result $\text{prefixsum}(s)$.
 - Go back to the original list, and again traverse from each s to $\text{next}(s)$ starting with the value $\text{prefixsum}(s)$ and adding the value at each node to a running sum and writing this into the node. Now all elements in the list should have the correct prefix sum.

Analyze the work and depth of this algorithm. These should both be given with high probability bounds (one idea might be to use Chernoff bounds for getting these).

2. **(6 points) Random Mate on Graphs** In class we described a random-mate technique for determining graph connectivity. Each node flips a coin, and every edge from a head to a tail will attempt to hook the tail into the head (i.e., relabel the tail with the head pointer). Given a d -regular graph on n vertices, i.e. a graph in which every vertex has degree d , what is the expected number of vertices after one contraction step?
3. **(8 points) Finding Parent Nodes In Tree:** Let $T = (V, E)$ be a rooted tree with $V = \{1, 2, \dots, n\}$. Each node in T has a specific ordering of its children. In other words, we have a specific first, second, ..., n -th children of a node. We can also define next sibling of i -th child of v as the $(i+1)$ -th child of v . You are given the value n along with the tree T as two arrays fc (first child) and ns (next sibling) of size n each. For each $u \in V$, the element $fc(u)$ is the first child of u , and 0 if u does not have children. Likewise, for each $u \in V$, the element $ns(u)$ is the next sibling of u , and 0 if u does not have a next sibling. The following figure gives an example for $n = 12$ and root node = 7, with left-to-right child ordering. We want to compute an array p of size n such that $p(u)$ stores the parent of u for all $u \in V$ (we will have $p(r) = 0$ for the root node r). Propose an $O(\log n)$ depth algorithm for an exclusive read exclusive write



PRAM model assuming exclusive reads have the same depth as concurrent reads for simplicity. Also specify what is the work done by your algorithm. You can use extra space, if required.

4. **(8 points) Parallel Gaussian Elimination:** Let A be an invertible $n \times n$ matrix with real-valued entries and we want to compute $B = A^{-1}$. One popular algorithm for this computation involves Gaussian elimination wherein we start with $B = I_n$ and apply elementary row operations (exchanging two rows, multiplying a row by a non-zero value, subtracting a multiple of a row from another) to convert A to the $n \times n$ identity matrix I_n . As we do this, we apply the same sequence of elementary row operations to B such that when A is reduced to I_n , B changes from I_n to A^{-1} . This is how the row-reduction algorithm looks like for doing operations on only A :

```

1 For i = 1, 2, ..., n, do:
2   - Find one k in {i, i+1, ..., n} such that A[k, i] != 0
3   - If k != i, swap row A[i, :] with row A[k, :]
4   - Divide row A[i, :] by A[i, i]
5   - For k in {1, 2, ..., n} with k != i and A[k, i] != 0:
6     Subtract A[k, i] times row A[i, :] from row A[k, :]

```

- (a) Design an $O(n \log n)$ depth parallel algorithm using the above matrix-inversion method and write its pseudocode. You can use extra space for your algorithm if required. Remember to mention the steps for initializing and computing entries in B as well when you use the above pseudocode as a scaffold for your solution.
 - (b) If it is given that A is a lower triangular invertible $n \times n$ matrix with real-valued entries, explain how we can develop a $O(\log^2 n)$ depth algorithm to invert A . Also deduce the work done by your algorithm. Can we gain anything on work complexity by using Strassen's algorithm? (Hint: Break each of A and A^{-1} into four $\frac{n}{2} \times \frac{n}{2}$ blocks, and use the fact that A^{-1} is again lower triangular.)
5. **(10 points) Stochastic Gradient Descent**

- (a) In class we stated that gradient descent on L -smooth functions is guaranteed to decrease the function value at each iteration. Stochastic gradient descent, on the other hand, does not have the same guarantee. Provide an example where stochastic gradient descent does not produce a descent step. Specifically, find a function $f(x) = \sum_{i=1}^m f_i(x)$, and an iterate x_0 such that for all step sizes, there exist i such that $f(x_1) > f(x_0)$ (where $x_1 := x_0 - \alpha \nabla f_i(x)$).

- (b) This exercise will guide you through the convergence proof of SGD. As a reminder, we are proving that if there exists a constant G such that $\mathbb{E}[\|\nabla f_i(x)\|^2] \leq G^2$ and $f(x)$ is μ -strongly convex. Then, with step-sizes $\gamma_k = \frac{1}{\mu k}$, we have

$$\mathbb{E}[\|x_k - x_*\|^2] \leq \frac{\max\{\|x_1 - x_*\|^2, \frac{G^2}{\mu^2}\}}{k}.$$

- Using strong convexity, prove that

$$\langle \nabla f(x_k) - \nabla f(x_*), x_k - x_* \rangle = \langle \nabla f(x_k), x_k - x_* \rangle \geq \mu \|x_k - x_*\|^2$$

- Apply the previous step, to express $\mathbb{E}[\|x_{k+1} - x_*\|^2]$ in terms of $\mathbb{E}[\|x_k - x_*\|^2]$, γ_k , G , and μ .
- Prove the convergence of SGD using induction.

6. (8 points) **HOGWILD!** Recall that in HOGWILD!, the objective function we want to minimize is :

$$f(x) = \sum_{e \in E} f_e(x_e)$$

where we define the hyperedge e to be the subset of variables that f_e depends on. Figure 1 depicts such a graph. Then, if we denote the average degree of the conflict graph as $\bar{\Delta}_C$, convergence is still guaranteed if the core delay is less than $\tau \leq \frac{n}{2\bar{\Delta}_C}$ (i.e., no more than τ samples are being processed while a core is processing one).

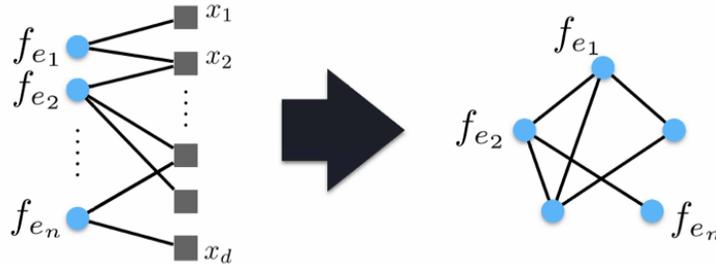


Figure 1: The function-variable and conflict graph for sparse functions.

Graph Cuts: In graph cuts problems, we are given a sparse matrix W which indexes similarity between node. We want to match each node to a list of D classes i.e., we want assign a vector $x_i \in \{v \in \mathbb{R}^D \mid \sum_{j=1}^D v_j = 1, v_j \geq 0\}$ that solve the following optimization problem.

$$\begin{aligned} & \underset{x}{\text{minimize}} && \sum_{(u,v) \in E} w_{uv} \|x_u - x_v\|_1 \\ & \text{subject to} && x_u \in \{v \in \mathbb{R}^D \mid \sum_{j=1}^D v_j = 1, v_j \geq 0\}. \end{aligned} \tag{1}$$

Prove that in this case

$$\frac{\bar{\Delta}_C}{n} = \mathcal{O}(\text{Avg. degree of graph in graph cut problem})$$

Programming Warm Up (30 points)

For the next few problems, we expect that you can learn SQL on your own and answer the below questions. This will act as a warm up for second half of the course. Some good guides for learning SQL:

- (a) https://www.w3schools.com/sql/sql_intro.asp
- (b) <https://www.youtube.com/watch?v=9Pzj7Aj25lw>

7. **(3 points)** In this set of problems, we review how to *select* data from relational databases.

- (a.) Write a SQL statement to find the total purchase amount of all orders. Sample table: `orders`.

<code>ord_no</code>	<code>purch_amt</code>	<code>ord_date</code>	<code>customer_id</code>	<code>salesman_id</code>
70001	150.5	2012-10-05	3005	5002
70009	270.65	2012-09-10	3001	5005
70002	65.26	2012-10-05	3002	5001
70004	110.5	2012-08-17	3009	5003
70007	948.5	2012-09-10	3005	5002

- (b.) Write a SQL statement which selects the highest grade for each of the cities of the customers. Sample table: `customer`.

<code>customer_id</code>	<code>cust_name</code>	<code>city</code>	<code>grade</code>	<code>salesman_id</code>
3002	Nick Rimando	New York	100	5001
3005	Graham Zusi	California	200	5002
3001	Brad Guzan	London		5005
3004	Fabian Johns	Paris	300	5006
3007	Brad Davis	New York	200	5001

- (c.) Write a SQL statement to find the highest purchase amount on a date “2012-08-17” for each salesman with their ID. Sample table: `orders`, used in (a).

8. **(8 points)** In this problem, we review how to *merge* two tables together.

- (a.) Write a SQL statement to know which salesman are working for which customer. Use the sample table `customer`, used in previous problem, and also `salesman`.

<code>salesman_id</code>	<code>name</code>	<code>city</code>	<code>commission</code>
5001	James Hoog	New York	0.15
5002	Nail Knite	Paris	0.13
5005	Pit Alex	London	0.11
5006	Mc Lyon	Paris	0.14
5003	Lauson Hen		0.12

- (b.) Write a query to display all salesmen and customers located in London.
- (c.) Write a SQL statement to make a cartesian product between salesman and customer i.e. each salesman will appear for all customer and vice versa for those salesmen who belongs to a city and the customers who must have a grade.
- (d.) Write a SQL statement to make a report with customer name, city, order number, order date, and order amount in ascending order according to the order date to find that either any of the existing customers have placed no order or placed one or more orders. Use customer and orders tables.

9. (5 points) In this problem, we consider *aggregation* of data.

- (a.) Write a SQL statement to find the highest purchase amount ordered by the each customer on a particular date with their ID, order date and highest purchase amount. Sample table: orders, used in problem 7 part (a).
- (b.) Write a SQL query to display the average price of each company's products, along with their code. Sample table: item_mast.

PROD_ID	PROD_NAME	PROD_PRICE	PROD_COMPANY
101	Mother Board	3200	15
102	Key Board	450	16
103	ZIP drive	250	14
104	Speaker	550	16
105	Monitor	5000	11
106	DVD drive	900	12

10. (4 points) **Joins with multiple keys** The point of this question is to explore how SQL handles cases where a join is performed on tables containing duplicate rows. Consider the following table item_mast.

PROD_ID	PRODUCT	PROD_PRICE	PROD_COMPANY
101	Mother Board	3200	1
101	Mother Board	2900	999
103	ZIP drive	250	14
106	DVD drive	900	12

and a corresponding table of customer purchases, purchases.

PROD_ID	CUSTOMER	PRODUCT	city
101	James Hoog	Mother Board	New York
101	James Hoog	ZIP drive	Los Angeles
103	Mc Lyon	ZIP drive	Pittsburgh

Notice that in `item_mast`, the same product can appear multiple times (listed under different manufacturers). Also, in database `purchases` the same customer can appear multiple times. If we join carefully using select columns, we can identify observations uniquely in the resulting output table. However, suppose we join the two tables only on `item_mast: product`, `product price` and `purchases: customer, product`.

Draw a sample table describing what the output looks like, and explain the result.

11. **(10 points)** Implement logistic regression using tensorflow. Use the following code to generate train and test data. Note that we have set seed (using "random_state=42"). Use cross-entropy loss and gradient descent optimizer with a learning rate of 0.01. Use `batch_size` of 100, and run for 500 steps. Report the accuracy on test set.

```
from sklearn.datasets import make_classification
from sklearn.model_selection import train_test_split
import matplotlib.pyplot as plt

# Generate data
X_data, y_data = make_classification(n_samples=200, n_features=2,
n_redundant=0, random_state=42)

# Split into train and test sets
X_train, X_test, y_train, y_test = train_test_split(X_data, y_data,
test_size=0.2, random_state=42)

# Plot training data
plt.scatter(X_train[:, 0], X_train[:, 1], c=y_train)
plt.show()
```